

APPARATUS AND METHOD FOR IMPLEMENTING RAID DEVICES IN A
CLUSTER COMPUTER SYSTEM

CROSS-REFERENCE TO RELATED APPLICATIONS

This application relates to: 6,728,905

10 U.S. Patent Application No. ~~097~~^{519,411}, entitled, "Apparatus and Method for Rebuilding a Logical Device in a Cluster Computer System," filed 3/3, 2000, naming Govindaraju Gnanasivam and Nallakumar Meiyappan as inventors, with Attorney Docket No. A-66978/JAS/RMA/LM / No. 38, and commonly assigned to Mylex Corp. of Fremont, CA;

15 U.S. Patent Application No. ~~097~~^{6,636,982}, entitled, "Apparatus and Method for Detecting the Reset of a Node in a Cluster Computer System," filed 3/3, 2000, naming Mohan B. Rowlands as inventor, with Attorney Docket No. A-66997/JAS/RMA/LM / No. 43, and commonly assigned to Mylex Corp. of Fremont, CA; and

20 U.S. Patent Application No. ~~097~~^{6,460,149}, entitled, "Suicide Among Well-Mannered Cluster Nodes Experiencing Heartbeat Failure," filed 4/11, 2000, naming Mohan B. Rowlands as inventor, with Attorney Docket No. A-67594/JAS/RMA/LM / No. 65, and commonly assigned to Mylex Corp. of Fremont, CA.

25

BACKGROUND

This invention relates to cluster computer systems in general. More particularly, the invention relates to recovering from cable failure in cluster computer systems with RAID devices.

30

power), that resource is simply added to the running cluster system, and the performance of the cluster system as a whole improves.

Well known in the art and only generally described here, the Microsoft Cluster Service product is the collection of all cluster-
5 activity-management software on each node of a Microsoft cluster system. The Cluster Service is more fully described in "Microsoft Windows NT Server Cluster Strategy: High Availability and Scalability with Industry-
Standard Hardware," (Microsoft Corp., 1995) and "Concepts and Planning: Microsoft 'Wolfpack' Clustering for Windows NT Server" (Microsoft Corp.,
10 1996). These two Microsoft clustering documents are attached hereto as Appendices A and B and are incorporated by reference as well.

A Microsoft cluster system uses the Small Computer Systems Interface (SCSI) bus with multiple initiators as the storage connection
(although Microsoft envisions supporting the Fiber Channel in the
15 future). Well known in the art, SCSI is an evolving standard directed toward the support of logical addressing of data blocks on data devices. Documents detailing the variations of SCSI over time (SCSI-1, SCSI-2 and SCSI-3, for example) are available from the American National Standards
Institute (ANSI) of New York, New York (www.ansi.org). SCSI-1, SCSI-2
20 and SCSI-3 are together referred to as "SCSI" herein.

Figure 1 illustrates a two-node cluster system 100
implemented on a SCSI bus 110 according to the prior art. In Figure 1, the cluster system 100 includes a first server node 120a and a second
server node 120b. The server nodes 120a and 120b have respective SCSI
25 identifiers (SCSI IDs) 7 and 6. The server nodes 120 connect to the SCSI bus 110 through respective host bus adapters (HBAs) 121.

A node 120 typically includes one or more of the following: a central processor unit ("CPU") 126, a memory 122, a user interface 123, a
co-processor 124, ports 125, a communications interface 121 and an
30 internal bus 127.

Of course, in an embedded system, some of these components may be missing, as is well understood in the art of embedded systems. In